

# Understanding Confounding Effects in Linguistic Coordination: an Information-Theoretic Approach

Shuyang Gao<sup>1,\*</sup>, Greg Ver Steeg<sup>1</sup>, Aram Galstyan<sup>1</sup>

**1 Information Sciences Institute, University of Southern California, Marina del Rey, CA**

**\* E-mail: Corresponding sgao@isi.edu**

## Abstract

We suggest an information-theoretic approach for measuring stylistic coordination in dialogues. The proposed measure has a simple predictive interpretation and can account for various confounding factors through proper conditioning. We revisit some of the previous studies that reported strong signatures of stylistic accommodation, and find that a significant part of the observed coordination can be attributed to a simple confounding effect - length coordination. Specifically, longer utterances tend to be followed by longer responses, which gives rise to spurious correlations in the other stylistic features. We propose a test to distinguish correlations in length due to contextual factors (topic of conversation, user verbosity, etc.) and turn-by-turn coordination. We also suggest a test to identify whether stylistic coordination persists even after accounting for length coordination and contextual factors.

## Introduction

Communication Accommodation Theory [1] states that people tend to adapt their communication style (voice, gestures, word choice, etc.) in response to the person with whom they interact. Originally, experiments on linguistic accommodation were confined to small scale laboratory settings with a handful of participants. The recent proliferation of digital (or digitized) communication data offers an opportunity to study nuances of human communication behavior on much larger scales. A number of recent studies have indicated presence of stylistic coordination in communication [2–5], where one person’s use of a linguistic feature (e.g. prepositions) increases the probability that a response will include the same feature. Linguistic style coordination (or *matching*) has been used to predict relationship stability [2] and negotiation outcomes [6], understand group cohesiveness [3], and infer relative social status and power relationships among individuals [5].

Most reports of linguistic style coordination have been based on correlational analysis. Thus, such claims are susceptible to various confounding effects. For instance, it is known that there is significant length coordination in dialogues, in the sense that a longer utterance from user  $Y$  tends to solicit a longer response from user  $X$  [7]. Thus, if the probability of an utterance containing a feature, e.g. *prepositions* or *words whose second letter is “r”*, depends only on length, this will create the illusion of stylistic coordination on the given feature.

Here we attempt to remedy the problem and propose an information-theoretic framework for characterizing stylistic coordination in dialogues. Namely, given a temporally ordered sequence of utterances (verbal or electronic statements depending on the context) by two individuals, we characterize their stylistic coordination with *time-shifted* mutual information. The proposed coordination measure characterizes the dependence between the stylistic features of the original post and the response. In addition, we provide a computational framework to account for confounders when measuring stylistic coordination.

We revisit some of the case studies where linguistic coordination was reported and demonstrate that a significant part of the observed correlations in linguistic features can indeed be explained by length coordination rather than stylistic accommodation. In particular, most stylistic features that exhibit statistically significant correlation exhibit little to no correlation after length coordination has been taken into account.

We also focus on the observed length correlations, and examine whether it is due to turn-by-turn coordination between the participants, or can be attributed to other contextual factors. We construct a statistical permutation test and demonstrate unequivocally that turn-by-turn length coordination in dialogues indeed takes place. Finally, we develop a similar test for stylistic features, and demonstrate that at least for one of the datasets, the remnant coordination (after conditioning on length) cannot be explained by contextual factors alone and has to be due to turn-by-turn level coordination between the speakers.

## Measuring Stylistic Coordination

### Representing Stylistic Features

To represent stylistic features in utterances, we use Linguistic Inquiry Word Count (LIWC) [8], which is a dictionary-based encoding scheme that has been used extensively for evaluating emotional and psychological dimensions in various text corpora. The latest version of the LIWC dictionary contains around 4500 words and word stems. Each word or word stem belongs to one or more word categories or sub-categories. Various LIWC categories include positive and negative emotion, function words, pronouns, articles, and so on. Here we focus on eight LIWC categories that have been used in previous studies [5]: articles, auxiliary verbs, conjunctions, high-frequency adverbs, impersonal pronouns, personal pronouns, prepositions, and quantifiers. Utterances are represented as eight-component binary vectors indicating the presence or absence of each linguistic marker [5].

### Information-theoretic measure of coordination

Each dialogue is a sequence of utterance exchanges between two participants. Following [4, 5] we binarize the stylistic features of utterances, so that a dialogue is represented as  $\{o_k^m, r_k^m\}_{k=1}^K$ , where  $o_k^m, r_k^m = \{0, 1\}$  indicate the absence or presence of the stylistic marker  $m$ , and  $K$  is the total number of exchanges in a dialogue. Since we focus on coordination between the same stylistic markers, we will drop the superscript  $m$  from now on. We use the convention  $O$  to represent the originator – the person who is producing the original utterance in a single exchange,  $R$  to represent the respondent – the person who is replying to the originator.

Let  $p(o, r)$  be the joint distribution of the random variables  $O$  and  $R$ . We characterize the amount of stylistic coordination using Mutual Information (MI) [9]; see S1 File for a brief overview of basic information theoretic concepts:

$$I(O : R) = H(O) - H(O|R) \quad (1)$$

where  $H(O) = -\sum p(o) \log p(o)$  is the Shannon entropy of  $O$ , and  $H(O|R)$  is the entropy of  $O$  conditioned on  $R$ . Note that in our case the arguments are temporarily ordered:  $O$  is always the initial utterance, and  $R$  is the response, so that Eq. 1 in fact defines time-shifted mutual information. Thus, even though MI is symmetric with respect to its argument, the coordination between two users may be asymmetric.

Recall that mutual information between two variables measures the average reduction in the uncertainty of one variable, if we know the other variable. Thus, in essence, the proposed measure of stylistic coordination quantifies how the use of a marker  $m$  in an utterance of  $O$ 's can help to predict  $R$ 's usage of  $m$  in the immediate response. In contrast to linear correlation measures, mutual information is well suited for handling strongly non-linear dependencies.

We measure the correlation between two variables *after* conditioning on a third variable,  $Z$ , via Conditional Mutual Information, defined as

$$I(O : R|Z) = H(O|Z) - H(O|R, Z). \quad (2)$$

Below we will use CMI to account for the confounding effect of the utterance length by conditioning on it. Namely, the actual stylistic accommodation, after accounting for the length coordination, is given by  $I(O : R|L_R)$ , where  $L_R$  is the length of the utterance by user  $R$ .

## Estimating mutual information from data

Given a set of samples  $\{o_k, r_k\}_{k=1}^K$ , our goal is to estimate mutual information between  $O$  and  $R$ . We could do this by first calculating the empirical distribution  $p(o, r)$  and then using Eq. 1. However, it is known that this naive *plug-in* estimator tends to underestimate the entropy of a system. Instead, here we use the statistical bootstrap method introduced by DeDeo et al. [10], which attempts to reduce the bias of the naive estimator by estimating a bootstrap correction term. The estimate of bias comes from comparing the entropy of the empirical distribution to estimates of entropy from several bootstrap datasets drawn randomly according to the empirical distribution. See [10] for more details.

While the above discrete estimator works well for evaluating mutual information between discrete stylistic variables, it is not very useful for evaluating mutual information between two length variables, due to limited number of samples we have. Instead, we will use a continuous estimator introduced by Kraskov et al. [11]. This non-parametric estimator searches the  $k$ -nearest neighbors to each point, and then average the mutual information estimated from the neighborhood of each point. It has been shown that this estimator is asymptotically unbiased and consistent. Discussion of different entropy estimators can be found in [12] and references therein.

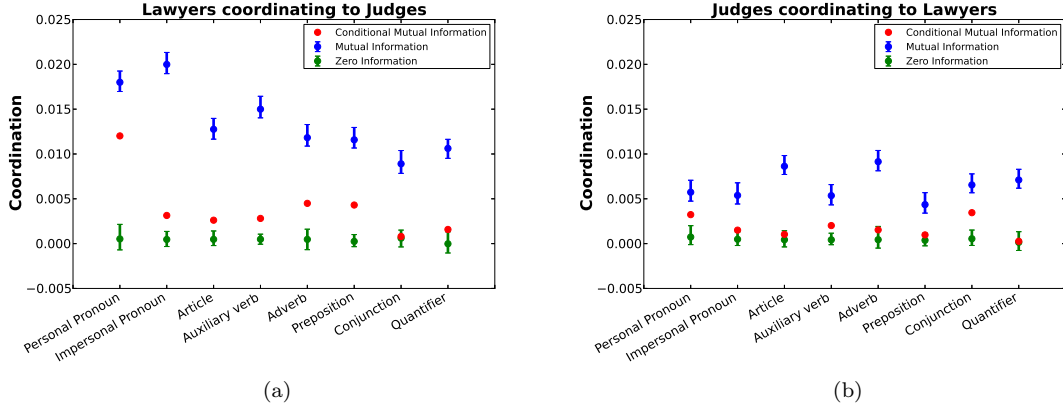
## Length as a confounding factor

We applied our coordination measures to two datasets previously studied in [5]: oral transcripts from the Supreme Court hearings, and discussion among Wikipedia editors. In the Supreme Court Data, there are 11 Judges and 311 Lawyers conversing with each other. We obtain 51,498 utterances from all the dialogues among 204 cases. In the Wikipedia dataset, users are classified into two categories, Administrators, or Admins, and *non-Admins*. All of the users interact with each other on Wikipedia *talk* pages, where they discuss issues about specific Wikipedia pages. We focus on dialogues where each participant make at least two exchanges within a dialogue, which results in over 30,000 utterances.

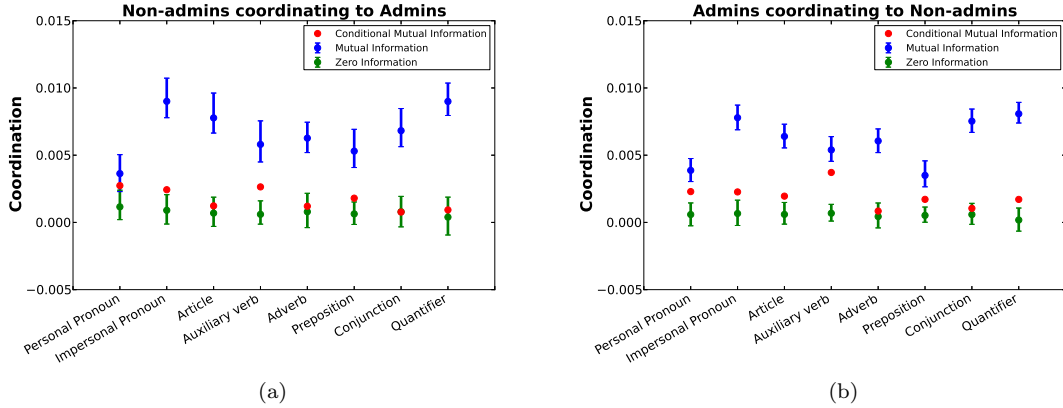
Ideally, we would like to calculate linguistic accommodation between any pair of individuals  $O$  and  $R$  who have participated in a dialogue. Unfortunately, most pair-wise exchanges are rather short and do not produce sufficient samples for evaluating mutual information or conditional mutual information. Instead, we group the individuals according to their roles, and then use aggregated samples to calculate stylistic coordination between the groups. The groups correspond to Judges and Lawyers for the Supreme Court data, and Admins and non-Admins for the Wikipedia data.

Fig. 1 describes stylistic coordination for the Supreme Court data as measured by  $I(O : R)$  and  $I(O : R|L_R)$ . The bias in estimators for conditional mutual information and mutual information are generally different. Therefore, rather than estimating mutual information directly, we use a conditional mutual information estimator where we condition on randomly permuted values for  $L_R$ . We repeat this procedure for four hundred times to produce 99% confidence intervals for  $I(O : R)$  (blue bars). The green bars give the 99% confidence intervals in case there is no stylistic coordination by estimating CMI with  $R$ 's utterances permuted (erasing any stylistic coordination).

The blue dots show the mutual information between the corresponding stylistic features, and suggest strong linguistic correlations between the groups. This effect, however, is strongly diminished after conditioning on the length of utterances (red dots). For instance, the coordination scores on features *Impersonal Pronoun*, *Article*, and *Auxiliary Verb* are reduced by factors of  $\sim 6.7$ ,  $\sim 4.8$ , and  $5.3$ , respectively, after conditioning on length. For the feature *Conjunction*, the 99% confidence interval of coordination score is above the confidence interval of *zero information* before conditioning, and falls into



**Figure 1. Coordination measures for the Supreme Court data.** The red (blue) dots give the true CMI (MI). The green dots represent CMI under the null hypothesis that there is no coordination after conditioning. (a) Lawyers coordinating to Judges. (b) Judges coordinating to Lawyers. In both figures, the conditional mutual information is significantly smaller than the mutual information for all eight stylistic features, indicating length is a confounding factor.



**Figure 2. Coordination measures for the Wikipedia data.** (a) Non-admins coordinating to Admins. (b) Admins coordinating to Non-admins. Symbols have the same interpretation as in the previous plot.

the confidence interval of *zero information* after conditioning. Similarly, in Fig. 1(b), the coordination scores for five out of eight markers (*Impersonal Pronoun*, *Article*, *Adverb*, *Preposition*, *Quantifier*) become practically zero after conditioning, suggesting that the observed coordination in those stylistic features are due to length correlations.

A similar picture holds for the Wikipedia dataset shown in Fig. 2. Again, we observe non-zero mutual information in all the features. However, this correlation is significantly diminished after conditioning on length. In fact, both non-admins coordinating to admins (Fig. 2(a)) and admins coordinating to non-admins (Fig. 2(b)) have an extremely weak signal after conditioning on length (all below 0.005). In particular, for non-admins coordinating to admins (Fig. 2(a)), the red dots of five out of eight features lie in the zero conditional mutual information confidence interval. For these five features in Fig. 2(b), we cannot rule out the null hypothesis that all stylistic coordination is due to phenomenon of length coordination.


Another interesting observation is that there is significant asymmetry, or *directionality*, in stylistic coordination. For instance, by comparing Figs. 1(a) and 1(b) we see that the mutual information is significantly higher from lawyers to judges than vice versa. A similar (albeit less pronounced) asymmetry is present for the Wikipedia data as well. This type of asymmetry has been used to suggest that the relative strength of stylistic accommodation reflects social status [5]. However, Figs. 1 and 2 illustrate that the asymmetry is drastically weakened after conditioning on length (red dots), suggesting that the phenomenon of higher stylistic coordination from lawyers to judges (and from non-admins to admins for the Wikipedia dataset) is due to the confounding effect of length. Unfortunately, a direct assessment of this effect in a single conversation is not feasible due to the insufficient number of utterances for calculating conditional mutual information. Nevertheless, in S2 File we suggest a different approach for addressing the above problem, and find that asymmetry in stylistic coordination can be explained by asymmetry in length coordination.

To conclude this section, we note that some of the correlations in stylistic features persist even after conditioning on length. One can ask whether this remnant correlation is due to turn-by-turn level linguistic coordination, or can be attributed to other confounding factors. We address this question in detail later in the text.

## Understanding Length Coordination

As discussed in the previous section, the observed correlations in linguistic features can be attributed to coordination in the length of utterances. Here we analyze this phenomenon in more detail. In particular, we are interested whether the observed length correlations are due to turn-by-turn coordination, or can be attributed to other contextual factors. For instance, consider a scenario that in one conversation, *Alice* and *Bob* are always conversing using short statements, while in another conversation they exclusively use long statements, perhaps due to different topics of conversation. Length coordination is found if data from these two conversation is aggregated, however, this coordination only reflects *Alice*’s and *Bob*’s response to the topic of conversation. More generally, aggregating data might lead to effects similar to Simpson’s paradox [13].

To understand the possible extent of various confounding factors (we call them *contextual factors*), consider the Bayesian network model that incorporates both *contextual factors* and *length coordination*, as shown in Fig. 3. Here  $L_O$  and  $L_R$  are random variables representing the length of an utterance by the originator  $O$  and the respondent  $R$ , respectively. In the model with both solid and dashed lines in Fig. 3(b),  $L_O$  explicitly influences  $L_R$ . While if we only have the solid lines in Fig. 3(a),  $L_R$  is independent of  $L_O$  after conditioning on the context  $C$ . Thus, the model in Fig. 3(a) assumes that there is only contextual coordination while Fig. 3(b) implies turn-by-turn coordination. Note that in principle, the contextual factor  $C$  can vary within a single conversation, for example, the theme of a conversation may change as time goes by. But for simplicity, we will assume that the contextual factor  $C$  does not



**Figure 3. A Bayesian network model for length coordination.** The network containing contextual factors,  $C$ , the length of an utterance,  $L_O^{(t)}$ , and the length of the response,  $L_R^{(t)}$ . (a) The lengths are correlated only due to contextual factors. (b) The lengths are correlated due to both contextual factors and potential effect of turn-by-turn level coordination (represented with the dotted line).

change within the dialogue or conversation.

### Information-theoretic characterization of length coordination

A direct measure of Turn-by-turn Length Coordination ( $TLC$ ) is given by the following conditional mutual information:

$$TLC = I(L_O : L_R | C) \quad (3)$$

Additionally, we define the Overall Length Coordination ( $OLC$ ) as

$$OLC = I(L_O : L_R) \quad (4)$$

Thus,  $OLC$  captures not only the length coordination in a turn-by-turn level, but also the confounding behaviors between  $L_O$ ,  $L_R$  and  $C$ . In fact,  $OLC$  can be decomposed into two items:

$$OLC = TLC + I(L_O : L_R : C) \quad (5)$$

The second item of right hand side in Eq. 5 indicates the multivariate mutual information (MMI) (also known as *interaction information* [14] or *co-information* [15]), and characterizes the amount of shared information between  $L_O$ ,  $L_R$  and  $C$ .

A straightforward method to test for turn-by-turn coordination is to evaluate  $TLC$  described in Eq. 3. Indeed,  $L_O$  and  $L_R$  are conditionally independent of  $C$  if and only if  $TLC = 0$ . However, direct evaluation of  $TLC$  is not possible due to the lack of sufficient number of samples, e.g., the number of exchanges within a specific dialogue. Nevertheless, it is possible to test the turn-by-turn length coordination by a non-parametric statistical test as shown below.

### Turn-by-Turn Length Coordination Test

Our null hypothesis is that there is no turn-by-turn coordination, so that all observed correlations are due to contextual factors. We now describe a procedure for testing this hypothesis.

We denote the pairwise set of exchanges in a specific dialogue  $c$  from originator  $o$  and respondent  $r$  as:

$$S_{o \leftarrow r}^c = \{o_c^k, r_c^k\}_{k=1}^{K_c} \quad (6)$$

where  $o_c^k, r_c^k$  indicate the  $k$ th exchange (two utterances) by the originator  $o$  and respondent  $r$  in dialogue  $c$ , and  $K_c$  represents the total number of exchanges in  $c$ . We also define the aggregated set of exchanges of user  $o \in O$  and user  $r \in R$  as:

$$S_{O \leftarrow R} = \bigcup_{o \in O, r \in R} \bigcup_{c \in C_{o,r}} S_{o \leftarrow r}^c \quad (7)$$

where  $C_{o,r}$  represents all the dialogues that involved user  $o$  and  $r$ . We can rewrite  $S_{O \leftarrow R}$  element-wise as

$$S_{O \leftarrow R} = \{O_k, R_k, C_k\}_{k=1}^N \quad (8)$$

where  $N = |S_{O \leftarrow R}|$  representing number of samples. For each triplet of right hand side in Eq. 8,  $R_k$  is the reply utterance to  $O_k$  in the dialogue  $C_k$ . Finally, from  $S_{O \leftarrow R}$  we obtain the set

$$L(S_{O \leftarrow R}) = \{\text{len}(O_k), \text{len}(R_k)\}_{k=1}^N \quad (9)$$

where  $\text{len}(\cdot)$  is a function representing the length of an utterance.

Consider now another sample, which is obtained by randomly permuting the respondent  $r$ 's utterances in the set  $S_{O \leftarrow r, c}$ :

$$\hat{S}_{o \leftarrow r}^c = \{o_c^k, \hat{r}_c^k\}_{k=1}^{K_c} \quad (10)$$

where  $\{\hat{r}_c^k\}_{k=1}^{K_c}$  is a random permutation of  $\{r_c^k\}_{k=1}^{K_c}$ . By aggregation, we have,

$$\hat{S}_{O \leftarrow R} = \bigcup_{o \in A, r \in B} \bigcup_{c \in C_{o,r}} \hat{S}_{o \leftarrow r}^c = \{O_k, \hat{R}_k, C_k\}_{k=1}^N$$


and

$$L(\hat{S}_{O \leftarrow R}) = \{\text{len}(O_k), \text{len}(\hat{R}_k)\}_{k=1}^N \quad (11)$$

Let us assume there is no turn-by-turn coordination, so that  $L_O$  and  $L_R$  are conditionally independent from each other given  $C$ . Then, it is easy to see that under this null model, the samples  $L(S_{O \leftarrow R})$  and  $L(\hat{S}_{O \leftarrow R})$  have the same likelihood, e.g., they are statistically equivalent. In other words,  $L(\hat{S}_{O \leftarrow R})$  can be viewed as a new sample from the same distribution  $p(l_o, l_r)$ . This observation suggests the following test: We first estimate  $OLC$  from the sample  $L(S_{O \leftarrow R})$  (denoted as  $OLC_0$ ) and then using the within-dialogue shuffled samples  $L(\hat{S}_{O \leftarrow R})$  (denoted as  $OLC_1$ ). Under the null hypothesis, these two estimates should coincide. Conversely, if  $OLC_0 \neq OLC_1$ , then the null hypothesis is rejected, suggesting that there is turn-by-turn length coordination.

The above procedure, which we call Turn-by-Turn Length Coordination Test, is a conditional Monte Carlo test [16]. The main advantage of this non-parametric test is that it requires a smaller sample size and does not need to make particular distribution assumptions. The test is non-parametric in two ways: the permutation procedure is non-parametric as well as the estimation of mutual information. We also note that in the context of stylistic coordination, a similar test was used in Ref. [17].

The results of this test are shown in Fig. 4. For the Supreme Court data, Fig. 4(a) shows that both Lawyers to Judges and Judges to Lawyers have non-zero mutual information ( $OLC_0$ ) before permutation. The Turn-by-Turn Length Coordination test shows that the mutual information decreases significantly after permutation (green confidence intervals,  $OLC_1$ ), rejecting the null hypothesis that  $L_O$  and  $L_R$  are



**Figure 4. Turn-by-turn length coordination test.** (a) Supreme Court dataset. (b) Wikipedia dataset. In both two subfigures,  $OLC_1$  is significantly smaller than  $OLC_0$ .

independent after conditioning on the contextual factor  $C$ . In other words, the contagion of length exists from the original utterance to the reply on a turn-by-turn level.

For the results on the Wikipedia discussion board in Fig. 4(b), we are also able to reject the null hypothesis. Notice that the degree of mutual information  $OLC$  is higher for Wikipedia than for the Supreme Court. However, one cannot make a general conclusion about the exact magnitude of turn-by-turn length coordination ( $TLC$ ) simply by calculating the loss, i.e.,  $OLC_0 - OLC_1$ .

## Revisiting Stylistic Coordination

We demonstrated in the previous section that strong correlations in utterance length explain most of the observed stylistic coordination. However, in some situations, there are statistically significant non-zero signals even after conditioning on length, e.g., the first feature dimension *Personal Pronoun* in Figs. 1(a) and 1(b). We now proceed to examine this remnant coordination. Specifically, we are interested in the following question: Does the non-zero conditional mutual information (after conditioning on length) represent turn-by-turn level stylistic coordination, or is it due to other contextual factors?

Toward this goal, consider the Bayesian network in Fig. 5, which depicts conditional independence relations between the length variables  $L_O$  and  $L_R$ ; stylistic variables  $F_O^m$  and  $F_R^m$  with respect to a style feature  $m$ , and the contextual (dialogue) variable  $C$ . The solid arrow from  $L_O$  to  $L_R$  reflects our findings from the last section about the existence of turn-by-turn length coordination. The dashed arc between the features  $F_O^m$  and  $F_R^m$  characterizes turn-by-turn stylistic coordination. Finally, the grey arcs between  $C$ ,  $F_O^m$  and  $C$ ,  $F_R^m$  indicate possible contextual coordination.


We use conditional mutual information to measure the Turn-by-turn Stylistic Coordination (TSC) with respect to a specific style feature  $m$ :

$$TSC = I(F_O^m : F_R^m | C, L_R) \quad (12)$$

where  $F_O^m$ ,  $F_R^m$  are binary variables indicating the feature  $m$  appears or not in an utterance. Also, the Overall Stylistic Coordination (OSC) is defined as

$$OSC = I(F_O^m ; F_R^m | L_R) \quad (13)$$





**Figure 5. A Bayesian network for linguistic style coordination.**  $L_O$  and  $L_R$  represent length of the respondent and length of the originator respectively.  $F_O^m$  and  $F_R^m$  represent a specific style feature variable for the respondent and originator.

Thus,  $OSC$  is exactly the conditional mutual information introduced in Eq. 2. Note that, even after conditioning on length,  $F_O^m$  and  $F_R^m$  are still dependent of each other because they are sharing the contextual factor  $C$ . ( $F_O^m \leftarrow C \rightarrow F_R^m$  is called a  $d$ -connected path in [18])

Again, a direct measure of turn-by-turn stylistic coordination corresponds to non-zero  $TSC$  in Eq. 12. However,  $TSC$  is hard to evaluate due to lack of sufficient samples. Furthermore, the shuffling test from the previous sections is not directly applicable here either, because it needs to be done in way that keeps the correlations between  $L_O$  and  $L_R$  intact: In other words, one can exchange utterances that have the same lengths. Since most dialogues are rather short, this type of shuffling test is not feasible, and one needs an alternative approach.

### Turn-by-Turn Stylistic Coordination Test

Our proposed test is based on the following idea: if we can rule out the influence of the contextual factors on stylistic correlations, then any non-zero conditional mutual information can be only explained by turn-by-turn stylistic coordination, i.e.,  $OSC = TSC$ . Thus, the null hypothesis is that there is contextual level coordination in stylistic features. We emphasize that by contextual coordination, we are actually referring to the links from  $C$  to  $F_O$  and  $C$  to  $F_R$  in Fig. 5.

We follow the same notation and methodology used in previous sections. By Eq. 8, let us denote the mixed length and stylistic feature set of  $S_{O \leftarrow R}$  as:

$$LF_m(S_{O \leftarrow R}) = \{len(O_k), len(R_k), f_m(O_k), f_m(R_k), C_k\}$$


where  $f_m(\cdot)$  is a binary function represents whether the style feature  $m$  in an utterance appears or not.

Consider now the shuffling procedure: we randomly permute respondent's utterances within a dialogue and obtain the set  $\hat{S}_{O \leftarrow R}$  in Eq. 11. We also define the length and feature set of  $\hat{S}_{O \leftarrow R}$ :

$$LF_m(\hat{S}_{O \leftarrow R}) = \{len(O_k), len(\hat{R}_k), f_m(O_k), f_m(\hat{R}_k), C_k\}$$

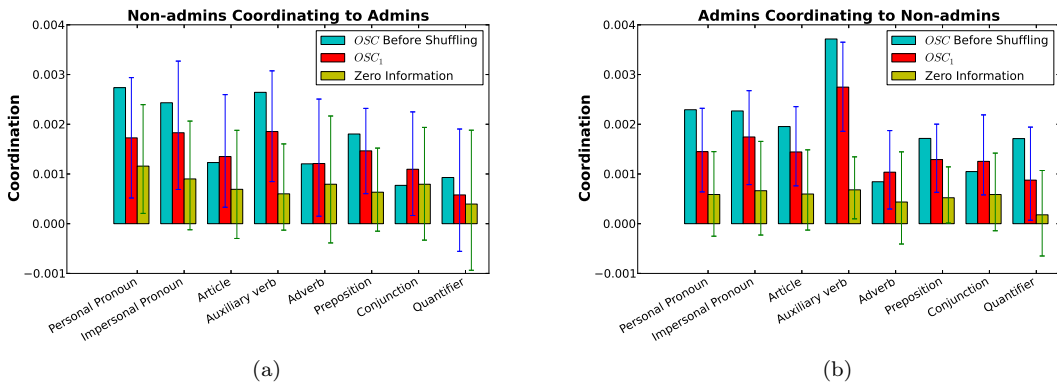
Clearly, the permutation destroys the turn-by-turn level coordination in both length and style. Thus, any remnant correlation must be due to contextual coordination, e.g., the fork  $F_O^m \leftarrow C \rightarrow F_R^m$ . This

provides a straightforward test for the existence of contextual coordination. Indeed, we simply need to estimate the overall stylistic coordination  $OSC_1$  using the shuffled sample  $LF_m(\hat{S}_{O \leftarrow R})$ . If  $OSC_1$  is larger than zero, then there is necessarily contextual coordination. On the other hand, if  $OSC_1 = 0$ , then all the observed stylistic correlations (calculated using the original non-shuffled sample) must be due to turn-by-turn stylistic coordination.



**Figure 6. Turn-by-turn stylistic coordination test for Supreme Court data.** (a) Lawyers coordinating to Judges. (b) Judges coordinating to Lawyers. (Blue bars indicate the overall stylistic coordination( $OSC$ ) before the test). One can see that after shuffling, values of  $OSC_1$  are within the zero-information confidence intervals.

Let us first consider the results of the above test for the Supreme Court data. From Figs. 6(a) and 6(b), one can see that for all the features, the corresponding CMI  $OSC_1$  are within the zero-information confidence intervals, indicating that non-zero conditional mutual information ( $OSC$  before shuffling) cannot be attributed to contextual factors. In other words, the remnant correlations that are not explained by length coordination must be due to turn-by-turn level coordination.



**Figure 7. Turn-by-turn stylistic coordination test for Wikipedia data.** (a) Non-admins coordinating to Admins. (b) Admins coordinating to Non-admins. (Blue bars indicate the overall stylistic coordination( $OSC$ ) before the test). One cannot rule out the null hypothesis that the remnant stylistic coordination is due to the contextual factors.

The situation is different for Wikipedia data. Indeed, Figs. 7(a) and 7(b) show that for the stylistic features with statistically significant remnant correlations even after conditioning on length ( $OSC$ ), the results of the above permutation tests are rather inconclusive. Namely, although the confidence intervals of  $OSC_1$  do overlap with the zero information confidence intervals, one cannot state unequivocally that they are zero. In other words, one cannot rule out the null hypothesis that the remnant stylistic coordination is due to the contextual factors rather than turn-by-turn coordination.

## Discussion

In conclusion, we have suggested an information theoretic framework for measuring and analyzing stylistic coordination in dialogues. We first extracted the stylistic features from the dialogue of two participants and then used *Mutual Information* ( $MI$ ) as a theoretically motivated measure of dependence to characterize the amount of stylistic coordination between the originator and the respondent in the dialogue. Moreover, by introducing *Conditional Mutual Information* ( $CMI$ ), which allows us to measure the correlation between two variables *after* conditioning on a third variable, we are able to more accurately gauge stylistic accommodation by controlling for confounding effects like length coordination.

We then used the proposed method to revisit some of the previous studies that had reported strong stylistic coordination. While the suggestion that one person’s use of, e.g., prepositions will (perhaps unconsciously) lead the other to use more prepositions is fascinating, our results indicate that previous studies have vastly overstated the extent of stylistic coordination. In particular, we showed that a significant part of the observed stylistic coordination can be attributed to the confounding effect of length coordination. We find that for both Supreme Court and Wikipedia data, the coordination score is greatly diminished after conditioning on length. We also find that the significant asymmetry in stylistic coordination shown in the previous study [5] is drastically weakened after conditioning on length. In fact, our results indicate that the asymmetry in length coordination can explain almost all the observed asymmetry in stylistic coordination.

Simpson’s paradox provides a famous example of how correlations observed in a population can disappear or even be reversed after conditioning on sub-populations. In an information-theoretic framework setting, a similar “paradox” can be seen in the example illustrated by Fig. 3: for  $L_O$ ,  $L_R$  and  $C$ , the mutual information  $I(L_O : L_R) > 0$ , while  $I(L_O : L_R|C) = 0$ . If we only look at the aggregated data, averaging over all contexts,  $C$ , i.e.,  $I(L_O : L_R)$ , there will be artificial mutual information between  $L_O$  and  $L_R$ . Ideally, we could calculate  $I(L_O : L_R|C)$  directly, however, there may not be enough samples for us to calculate the conditional mutual information for all values of  $C$ . How can we still determine whether  $I(L_O : L_R|C)$  is zero or not while using all the data? We thus designed non-parametric statistical tests to solve this problem in general while making full use of the available data. More importantly, because these information-theoretic quantities directly reflect constraints on graphical models, the mystery of Simpson-like paradoxes is replaced with concrete alternatives for generative stories as depicted in Figs. 3 and 5.

We also observed that for some of the stylistic markers, there was diminished but still statistically significant correlations even after conditioning on length. We again designed a non-parametric statistical test for analyzing this remnant coordination more thoroughly. Our findings suggest that for the Supreme Court data, the remnant coordination cannot be fully explained by other contextual factors. Instead, we postulate that the remnant correlations in the Supreme Court data is due to turn-by-turn level coordination. For the Wikipedia data, however, our results are less conclusive, and we cannot draw any conclusion about turn-by-turn stylistic coordination. Thus, caution must be taken when making general claims about the possible origin of stylistic coordination in different settings.

It is possible to develop alternative tests based on a more fine-grained, token-level generative models. The main idea behind such a test is to shuffle the word tokens uttered by an individual within each dialogue, which should destroy turn-by-turn coordination. Our preliminary results based on this test

suggest that most of the remnant correlations are indeed due to turn-by-turn coordination. However, we emphasize that this test requires an additional assumption whose validity needs to be verified, namely that the words used by a given speaker within a conversation are independent and identically distributed (i.i.d.). Furthermore, the test assumes stationarity, i.e., that the contextual factors do not vary within the course of the dialogue. While this assumption seems reasonable in the dialogue settings considered here, it is important to note that deviations from stationarity might be yet another serious obstacle for identifying stylistic influences [19, 20]. Indeed, if we relax the stationarity condition, then any observed correlation in stylistic features might be due to temporal evolution rather than direct influence. And since any permutation-based test destroys temporal ordering, it cannot differentiate between those two possibilities.

While our work focuses on linguistic style matching, we believe that the information-theoretic method proposed can be useful for studying more general types of linguistic coordination in dialogues, such as structural priming [21, 22], or lexical entrainment [23, 24]. Recall that according to the structural priming hypothesis, the presence of a certain linguistic structure in an utterance affects the probability of seeing the same structure later in the dialogue. This type of turn-by-turn coordination can naturally be captured by (time-shifted) mutual information between properly defined linguistic variables. Furthermore, using the permutation tests described here, it should be possible to differentiate between *historical* and *ahistorical* mechanisms of lexical entrainment [24]. Indeed, the former mechanism assumes some type of influence/coordination between the speakers that helps them to arrive at a common conceptualization. The ahistorical mechanism, on the other hand, assumes that the speaker’s choice of each term is an independent event affected by the informativeness and availability of the term, and some other factors, which, in our terminology, is analogous to contextual coordination.

In a broader context, we note that sociolinguistic analysis has been used for assessing and predicting societally important outcomes such as health behaviors, suicidal intent, and emotional well-being, to name a few examples [25–29]. Thus, it is imperative that such predictions are based on sound theoretical and methodological principles. Here we suggest that information theory provides a powerful computational framework for testing various hypotheses, and furthermore, is flexible enough to account for various confounding variables. Recent advances in information-theoretic estimation are shifting these approaches from the theoretical realm into practical and useful techniques for data analysis. We hope that this work will contribute to the development of mathematically principled tools that enable computational social scientists to draw meaningful conclusions from socio-linguistic phenomena.

## Acknowledgments

This research was supported in part by Defense Advanced Research Projects Agency (DARPA) grant No. W911NF-12-1-0034.

## References

1. Giles H, Coupland N, Coupland J (1991) Accommodation theory: Communication, context, and consequence, Cambridge: Cambridge University Press. pp. 1–68.
2. Ireland ME, Slatcher RB, Eastwick PW, Scissors LE, Finkel EJ, et al. (2011) Language style matching predicts relationship initiation and stability. *Psychological Science* 22: 39–44.
3. Gonzales AL, Hancock JT, Pennebaker JW (2010) Language style matching as a predictor of social dynamics in small groups. *Communication Research* 37: 3–19.
4. Danescu-Niculescu-Mizil C, Gamon M, Dumais S (2011) Mark my words!: linguistic style accommodation in social media. In: *Proceedings of the 20th international conference on World wide*

- web. WWW '11, pp. 745–754. doi:10.1145/1963405.1963509. URL <http://doi.acm.org/10.1145/1963405.1963509>.
5. Danescu-Niculescu-Mizil C, Lee L, Pang B, Kleinberg J (2012) Echoes of power: language effects and power differences in social interaction. In: Proceedings of the 21th international conference on World wide web. WWW '12, pp. 699–708.
  6. Taylor PJ, Thomas S (2008) Linguistic style matching and negotiation outcome. *Negotiation and Conflict Management Research* 1: 263–281.
  7. Niederhoffer KG, Pennebaker JW (2002) Linguistic style matching in social interaction. *Journal of Language and Social Psychology* 21: 337–360.
  8. Pennebaker JW, Francis ME (2007). Linguistic inquiry and word count: A computerized text analysis program. URL <http://www.liwc.net/>.
  9. Cover TM, Thomas JA (2006) *Elements of Information Theory* (Wiley Series in Telecommunications and Signal Processing). Wiley-Interscience.
  10. DeDeo S, Hawkins RXD, Klingenstein S, Hitchcock T (2013) Bootstrap methods for the empirical study of decision-making and information flows in social systems. *Entropy* 15: 2246–2276.
  11. Kraskov A, Stögbauer H, Grassberger P (2004) Estimating mutual information. *Physical Review E* 69: 066138.
  12. Wang Q, Kulkarni SR, Verdú S (2009) Universal estimation of information measures for analog sources. *Found Trends Commun Inf Theory* 5: 265–353.
  13. Simpson EH (1951) The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society, Series B (Statistical Methodology)* 13: 238–241.
  14. McGill W (1954) Multivariate information transmission. *Information Theory, Transactions of the IRE Professional Group on* 4: 93–111.
  15. Bell AJ (2003) The co-information lattice. In: Proceedings of the 4th international symposium on Independent Component Analysis and Blind Source Separation. pp. 921–926.
  16. Pesarin F, Salmaso L (2010) *Permutation Tests for Complex Data: Theory, Applications and Software*. John Wiley & Sons.
  17. Danescu-Niculescu-Mizil C, Lee L (2011) Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs. In: Proceedings of the 2Nd Workshop on Cognitive Modeling and Computational Linguistics. Stroudsburg, PA, USA: Association for Computational Linguistics, CMCL '11, pp. 76–87. URL <http://dl.acm.org/citation.cfm?id=2021096.2021105>.
  18. Pearl J (2000) *Causality: Models, Reasoning, and Inference*. New York, NY, USA: Cambridge University Press.
  19. Ver Steeg G, Galstyan A (2013) Minimal assumption tests for contagion in observational social network studies. In: Workshop on Information in Networks(WIN).
  20. Ver Steeg G, Galstyan A (2013) Statistical tests for contagion in observational social network studies. In: AISTATS'13.
  21. Bock JK (1986) Syntactic persistence in language production. *Cognitive Psychology* 18: 355–387.

22. Pickering MJ, Ferreira VS (2008) Structural priming: a critical review. *Psychological bulletin* 134: 427.
23. Brennan SE (1996) Lexical entrainment in spontaneous dialog. In: *International Symposium on Spoken Dialog*. pp. 41–44. URL [citeseer.ist.psu.edu/brennan96lexical.html](http://citeseer.ist.psu.edu/brennan96lexical.html).
24. Brennan SE, Clark HH (1996) Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22: 1482.
25. Pennebaker JW, King LA (1999) Linguistic styles: Language use as an individual difference. *Journal of Personality and Social Psychology* 77: 1296–1312.
26. Stirman SW, Pennebaker JW (2001) Word use in the poetry of suicidal and nonsuicidal poets. *Psychosomatic Medicine* 63: 517–522.
27. Campbell RS, Pennebaker JW (2003) The secret life of pronouns. *Psychological Science* 14: 60–65.
28. Rude S, Gortner EM, Pennebaker J (2004) Language use of depressed and depression-vulnerable college students. *Cognition and Emotion* 18: 1121+.
29. Boals A, Klein K (2005) Word use in emotional narratives about failed romantic relationships and subsequent mental health. *Journal of Language and Social Psychology* 24: 252–268.

## Supporting Information

**S1 File. Basic Concepts in Information Theory.**

**S2 File. Stylistic Coordination and Power Relationship.**

## S1 Basic Concepts in Information Theory

Consider a random variable  $X$  with probability distribution  $p(x) \equiv p(X = x)$ . In the discrete case, the *Shannon* entropy is defined as:

$$H(X) = - \sum_x p(x) \log p(x) \quad (14)$$

Note that If  $X$  is a continuous variable, the sum in Eq. 14 is replaced by an integral. We often talk about the entropy  $H(X)$  as quantifying our uncertainty about the random variable, so that higher entropy means more uncertain (and less predictable)  $X$ . In particular, entropy is zero,  $H(X) = 0$ , if and only if  $X$  is perfectly predictable.

Let us now consider two random variables,  $X$  and  $Y$ , and let  $p(x, y)$  denote their joint distribution. The (joint) entropy for  $X$  and  $Y$  is defined similarly as

$$H(X, Y) = - \sum_x \sum_y p(x, y) \log p(x, y) \quad (15)$$

We can also define conditional entropy of  $Y$  given  $X$  (or vice versa), as follows:

$$H(Y|X) = - \sum_x p(x) \sum_y p(y|x) \log p(y|x) \quad (16)$$

If  $X$  and  $Y$  are independent, then their joint entropy is the sum of individual entropies,  $H(X, Y) = H(X) + H(Y)$ . Indeed, note that for independent  $X$  and  $Y$ , their joint distribution factorizes as  $p(x, y) = p(x)p(y)$ . Plugging this factorized distribution into Eq. 15 then yields the additivity of the joint entropy.

In case  $X$  and  $Y$  are not independent, the degree of their dependence can be measured by *mutual information* (MI), defined as follows:

$$I(X : Y) = H(X) + H(Y) - H(X, Y) \quad (17)$$


$I(X : Y)$  measures the correlation between  $X$  and  $Y$ , and equals to zero if and only if  $X$  and  $Y$  are independent. In essence,  $I(X : Y)$  measures the amount of uncertainty reduction in  $X$ , if we know  $Y$ , and vice versa. This intuition becomes more clear if we rewrite the mutual information in the following form:

$$I(X : Y) = H(Y) - H(Y|X) \quad (18)$$

$$= H(X) - H(X|Y) \quad (19)$$

The relationships between the entropies, conditional entropies, and mutual information are captured in the Venn diagram shown in Fig. A.

For a more detailed account of information-theoretic concepts, we refer the reader to a classical textbook by Cover and Thomas [9].



**Figure A.** The Venn diagram depicting the relationship between individual ( $H(X)$ ,  $H(Y)$ ), joint ( $H(X, Y)$ ), and conditional ( $H(X|Y)$ ,  $H(Y|X)$ ) entropies. The intersection of the circles is the mutual information  $I(X; Y)$ .



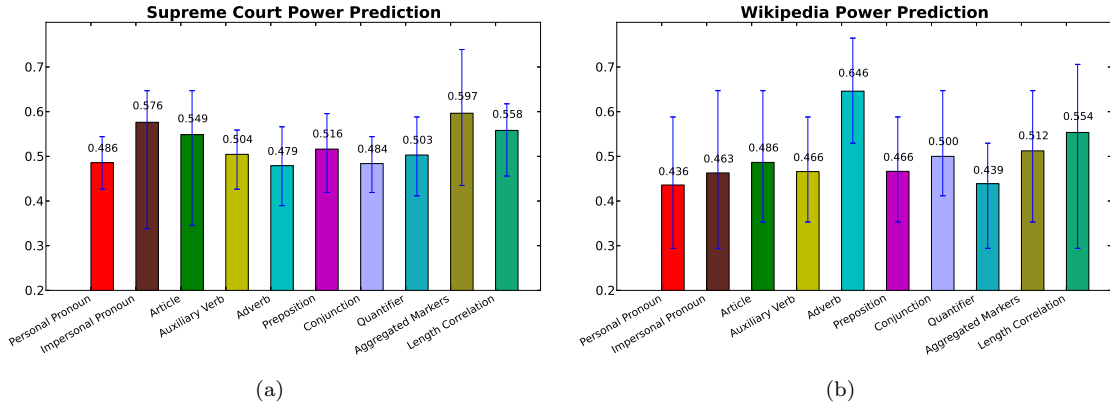
## S2 Stylistic Coordination and Power Relationship

It has been hypothesized that directionality of the stylistic coordination in dialogues can be predictive of power relationship between the conversations, as discussed in [5]. We do indeed observe directional differences in stylistic coordination when comparing Figs. 1(a) and 1(b) with Figs. 2(a) and 2(b). However, as we elaborated above, the observed directionality can result from the confounding effect of length coordination.

Here we analyze this issue in more details by setting up the following prediction task (see [5]). We consider all the pairs of users ( $X, Y$ ) who have different social status (e.g., admin vs. non-admin) and have engaged in dialogues. We then calculate stylistic coordination scores from  $X$  to  $Y$  and  $Y$  to  $X$ , and examine whether those scores can be used to classify the social status of each speaker. For classification, we assume we know the status relationship for a fraction of pairs in our dataset, and then use a supervised learning method called Support Vector Machine (SVM) to predict the status of the unknown users. We perform this prediction tasks using the following three different set of features:

- *Coordination Features*: For each pair, and for each of the eight stylistic markers, we produce two-dimensional feature vector, where the two components correspond to the mutual information score in either direction.
- *Aggregated Coordination Features*: For each pair, we aggregate the *Coordination Features* for all eight stylistic markers in both directions, which results in a sixteen-dimensional feature vector.
- *Length Coordination Features*: For each pair, we calculate the Pearson correlation coefficients between length of utterances in either direction, and use those coefficients as an input to SVM.

For the Wikipedia data, we consider (admin, non-admin) pairs and for the Supreme Court data, we consider (justice, lawyer) pairs.



**Figure A.** SVM Prediction Accuracy for both stylistic coordination features and length coordination features

In our experiment, we only select pairs which have at least 20 exchanges between them, so that we can calculate mutual information with reasonable accuracy. This results in 135 pairs in Supreme Court dataset and 34 pairs in Wikipedia dataset. Also, we labeled half fraction of the pairs, shuffled the training data and repeated the procedure  $N = 100$  times to calculate the average prediction accuracy.

Fig. A depicts the prediction accuracy for each of the above scenario, together with error bars which give the 95% confidence intervals. Since the two datasets are small, the error bars are relatively large in

these situations. The results can be summarized as follows. For the Supreme Court dataset, the best prediction accuracy is achieved when using *Aggregated Coordination Features*, whereas for the Wikipedia dataset the best accuracy corresponds to using coordination on the feature *Adverb*. More importantly, we find that using *Length Coordination Features* alone can predict user status with better-than-random (50%) accuracy. In fact, investigating from the error bars, we cannot rule out that the hypothesized correlation between social status and (the direction of) stylistic coordination is due to the confounding effect of length coordination.